

Goal-Oriented Regulatory Intelligence: How Can Watson Analytics Help?

Okhaide Akhigbe^(✉), Susie Heap, Sakib Islam, Daniel Amyot,
and John Mylopoulos

School of Electrical Engineering and Computer Science,
University of Ottawa, Ottawa, Canada
{okhaide, sheap069, sislao62, damyot,
jmylopou}@uottawa.ca

Abstract. Regulations are introduced by governments to ensure the well-being, safety, and other societal needs of citizens and enterprises. Governments also create programs aiming to improve awareness about and compliance with regulations. Goal models have been used in the past to conceptualize regulations and to measure compliance assessments. However, regulators often have difficulties assessing the performance of their regulations and programs. In this paper, we model both regulations and regulatory programs with the Goal-oriented Requirement Language. Using the same conceptualization framework enables asking questions about performance and about the evidence-based impact of programs on regulations. We also investigate how Watson Analytics, a cloud-based data exploration service from IBM, can be used pragmatically to explore and visualize goal satisfaction data to understand compliance issues and program effectiveness. A simplified example inspired from a Canadian mining regulation is used to illustrate the many opportunities of Watson Analytics in that context, and some of its current limitations.

Keywords: Data analytics · Data visualization · Goal models · Goal-oriented Requirement Language · GoRIM · Regulatory compliance · Regulatory intelligence · Watson Analytics

1 Introduction

Based on government policy objectives, *regulations* are introduced to ensure the well-being and safety of citizens and enterprises. Regulations aim to constrain behaviors of citizens and enterprises alike to achieve desired societal outcomes [1]. Governments also introduce and manage regulatory *programs*, which consist of events, items, activities, or processes for ensuring compliance to regulations. Regulatory programs improve awareness about and compliance with regulations by educating regulated parties about obligations and rights in relation to a regulation, and by promoting and monitoring compliance through inspections and other means [2]. While regulations routinely evolve throughout their lifetime to ensure they continue to address societal needs, it is often unclear whether they actually achieve intended societal outcomes. Do regulatory programs result in improved compliance? What do observed

compliance levels tell us about a regulation or its supporting programs? Are they meeting their intended objectives? Over the years, governments, citizens, and interest groups have been actively involved in answering such questions, often with inconclusive results [3–5]. Current trends, including climate change and cyber threats, are driving an increased interest to quantitatively link programs and regulations with societal outcomes.

Goal models have been used successfully in the past to conceptualize and analyze regulations. Goal models capture the structure and intent of regulations, and enable compliance measurements and assessments [6, 7]. In this context, goal models put different compliance scenarios in proper perspective for stakeholders to visualize relevant regulations, laws, processes, and objectives. If we could model regulatory programs in the same manner, we could then exploit a uniform modelling framework to ask questions about performance and the evidence-based impact of programs on regulations. Such an approach will support the monitoring, analysis, and assessment of regulations and their supporting regulatory programs. Regulators collect and use much data while administering (i.e., introducing, enforcing, reviewing, and evolving) regulations. Judging from the numerous regulated parties a regulation can influence, these data exhibit the three V-properties of Big Data (velocity, variety, volume) [8]. Goal models used in regulatory contexts collect data from numerous sources that also have velocity, variety, and volume. Visualizing and deriving insight from such Big Data is very challenging today because of the dimensions required for proper analysis, including conventional ones such as time, location, and organizations, but also domain-specific ones such as the structure of regulations and programs. Towards this end, we investigate the use of *IBM Watson Analytics* [9], a cloud-based Big Data technology, to explore and visualize these different dimensions of data in order to understand compliance issues and program effectiveness at the heart of many challenges faced by regulators [10].

One contribution of this article is our proposal to use the same conceptualization framework to model both regulations and their supporting regulatory programs. This approach enables us to obtain homogeneous goal satisfaction data from the goal models of regulations and regulatory programs. The other contribution of this paper is an extended method for regulatory intelligence that exploits *Watson Analytics* to explore and visualize the evaluated goal satisfaction data to obtain useful insight on the regulatory process. Here, the dimensions of data analysis are the structure of the regulation and program, location, and time. We demonstrate our method using an illustrative case study inspired from the Canadian mining sector to show the potential of this tool-supported conceptualization for supporting and enhancing regulatory practices.

The rest of the paper is as follows. Section 2 provides background on the use of goal models for regulatory compliance. Section 3 discusses regulatory intelligence and how it relates to the regulatory ecosystem. Section 4 introduces the Goal-oriented Regulatory Intelligence Method, while Sect. 5 uses an illustrative case study inspired from a real regulation to describe its applicability. We present lessons learned in Sect. 6 and limitations in Sect. 7. We conclude with a summary and future work in Sect. 8.

2 Background

Goal models are often used to show compliance of information systems and business processes with one or more regulations. Here, goal models exploit various concepts (such as goals, links, and actors) to assess compliance and explore what-if scenarios to address non-compliance. The rationale is that if goal models are a useful conceptualization for eliciting, modeling, and analyzing requirements in order to capture alternatives and conflicts between stakeholder objectives [6], they can also help explore and analyze compliance [7]. The Nòmos framework [11], including its variations (Nòmos 2 [12] and Nòmos 3 [13]), is a goal-based modeling framework used to systematically generate law-compliant requirements and support requirements analysts in dealing with the problem of requirements compliance. Secure Tropos [14] is another goal-based conceptualization that has been used to support the consideration of laws and regulations during the development of secure software systems. Finally, the User Requirements Notation (URN) [15], a standard modeling language used to model and analyze requirements with two complementary views, namely the *Goal-oriented Requirement Language* (GRL) and the Use Case Map (UCM) notation, has been used to model and study the compliance of enterprises goals and business processes against regulations [16, 17]. As of 2012, the URN standard includes an indicator concept that enables enhanced compliance analysis [18, 19], and regulator-oriented reasoning about the suitability of regulations and opportunities for their evolution [20, 21].

However, these goal modelling approaches have taken into consideration neither regulatory programs nor the Big Data aspects involved in the regulatory process. Administering regulations involves more than enforcing compliance. As such, regulators, citizens, and enterprises need to get insight from the data involved in the regulatory process. Data analytics technologies such as Watson Analytics can facilitate this. Watson Analytics is a pioneering software system that uses cloud computing and multiple machine learning algorithms to analyze high volumes of data [9, 22]. Using a simple intuitive user interface, Watson Analytics enables the user to ask questions on the collected data in natural language and returns results mined from the data across different dimensions of interest. Watson Analytics understands complicated and difficult questions asked in natural language, gives evidence-based results in an appropriate visualization, and proposes related questions of potential interest about patterns, trends, and correlations. There is growing acceptance and use of Watson Analytics, and some companies have recently started using it in a regulatory compliance context [23]. However, to our knowledge, Watson Analytics has not been used from a regulator's perspective, nor has it been used with goal models, until now.

3 Regulatory Intelligence

The concept of regulatory intelligence has its origins in the heavily regulated pharmaceutical industry [24]. The motivation for regulatory intelligence is to enable pharmaceutical companies to remain locally and globally compliant to existing and new regulations. As such, definitions of *regulatory intelligence* revolve around continuously obtaining and processing data and information from multiple sources and

analyzing them in the relevant context. It also includes generating and communicating meaningful outputs from these data in line with an organization's regulatory strategy [24, 25]. This implies that with regulatory intelligence, information relating to a given compliance context and its implication can be obtained, analyzed, and communicated. The regulatory ecosystem is also monitored to identify opportunities where insight obtained from the collected information can be utilized to influence future regulations. This application of regulatory intelligence is to improve decisions making and planning for pharmaceutical companies. It enables them to make and maintain their products compliance with regulations [26].

Beyond the pharmaceutical industry, in other domains, the interactions between regulators, citizens, and enterprises already involve some sort of data gathering, analysis, and communication about regulations and the regulatory process. Hence, regulatory intelligence is conducted *by regulators* using feedback from citizens and enterprises, and compliance enforcement information to administer regulations [21]. As such, from the regulator's perspective, regulatory intelligence can be used to enhance the regulatory process with data-driven support for decision-making towards introducing, enforcing, reviewing, and evolving regulations. Regulatory intelligence facilitates monitoring and assessing regulations and can be used to influence the regulatory process and ecosystem. Furthermore, analyzed data and information can be used to ascertain the relevance, effectiveness (e.g., in terms of goal satisfaction), or efficiency (e.g., in terms of costs/benefits) of regulations and their supporting regulatory programs. The regulator-oriented view of regulatory intelligence is the one adopted in this article.

4 A Method for Regulatory Intelligence

The concept of regulatory intelligence alludes to a feedback loop in the use of data from and within the regulatory ecosystem to administer regulations. In 2013, Badreddin et al. [21] proposed a regulatory intelligence method based on GRL that enables reasoning about regulations and compliance with regulation as a dimension. As discussed in Sect. 2, this method did not take regulatory programs and the amount of Big Data involved in the regulatory process, into consideration. We extend this method by incorporating a step that exploits Watson Analytics to provide a pragmatic way to explore and visualize regulations and regulatory programs as dimensions for data analysis. In addition, we explore the use of Watson Analytics to analyze the Big Data resulting from the evaluated goal models of regulations and regulatory programs to gain insights about the regulatory process. Our proposal, the *Goal-oriented Regulatory Intelligence Method* (GoRIM), shown in Fig. 1, is inspired from the method introduced by Tawhid et al. [20] for managing outcome-based regulations.

As a starting point, in the first step (**Build**), GRL models of the regulation and of the regulatory program are built using *jUCMNav*, a free Eclipse-based plugin for URN modeling and analysis [19, 20]. These models are built using the semi-automatic method for creating goal models of regulations from tables described by Rashidi-Tabrizi et al. [27]. The same GRL concepts are used for both types of models: goals, indicators, contribution/decomposition links, actors (optional), and dependencies

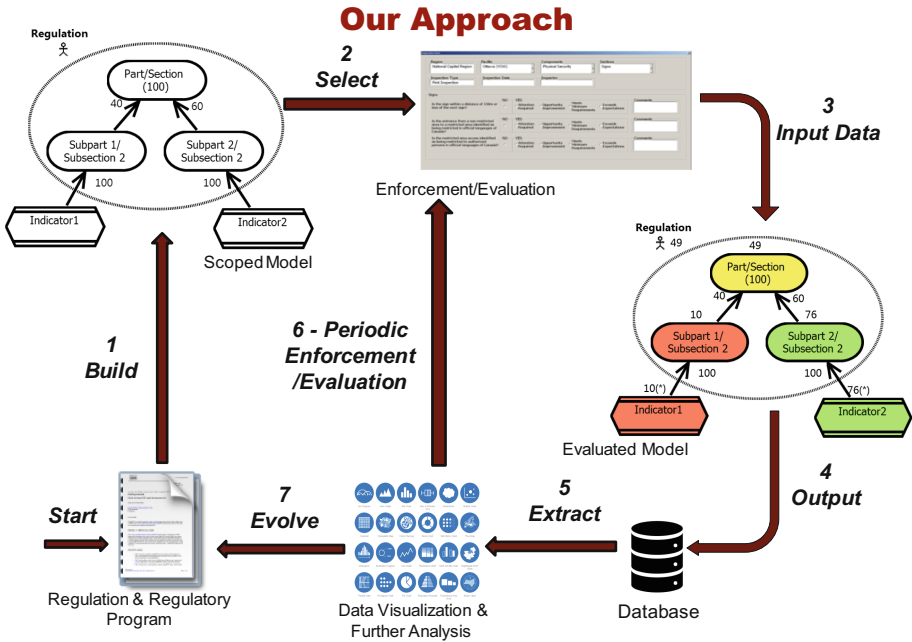


Fig. 1. The Goal-oriented Regulatory Intelligence Method (GoRIM)

to resources for conditional parts of regulations/programs (optional). Contribution levels and indicators are added manually by experts to the tabular representation of the regulations/programs as they are typically not found in the original documents. jUCMNav creates goal models by importing the tabular representation (comma-separated value file created with Excel), a format commonly used by regulators [27].

In the second step (*Select*), questions to be answered by inspectors/auditors (or regulated parties themselves in case of self-reporting) during periodic compliance enforcement activities for the regulation, as well as evaluations of the regulatory programs, are selected from predefined questions so that data can be fed to the indicators in the goal models. In the third step (*Input Data*), the data collected are input to the goal models as GRL strategies. Using GRL evaluation algorithms [28], satisfaction levels for both the regulation and program goal models, which indicate compliance and performance levels, are computed for all goals. In the fourth step (*Output*), snapshots of different computed compliance and performance levels can be produced for different regulated parties (companies, provinces, etc.) at different times, and stored in a database. In the fifth step (*Extract*), the data is extracted from the database and input into a data visualization engine (such as Watson Analytics in our case). Visualizations and further analysis can be done on large datasets to enable reporting on computed compliance and performance levels and what they mean relative to the regulation and regulatory program. Based on these computed levels, the needs for reinforcements or reevaluations can be highlighted in the sixth step (*Periodic Enforcement/Evaluation*). Decisions can be made on specifics to focus

on during the next rounds of enforcement or evaluations. In the seventh step (Evolve), the needs for evolution (addition, change, or repeal) of the regulation and/or program can be triggered based on the insight gained in the fifth step.

5 Illustrative Case Study

To illustrate GoRIM, we apply it to an example inspired from the Canadian *Metal Mining and Effluent Regulation* (MMER) [29]. A fragment of MMER is shown in Fig. 2.

<small>Metal Mining Effluent Regulations</small> <small>PART 1 General</small> <small>Authority to Deposit in Tailings Impoundment Areas</small> <small>Sections 5-7</small>	<small>Règlement sur les effluents des mines de métaux</small> <small>PARTIE I Dispositions générales</small> <small>Autorisation de rejeter dans un dépôt de résidus miniers</small> <small>Articles 5-7</small>
<p>Authority to Deposit in Tailings Impoundment Areas</p>	<p>Autorisation de rejeter dans un dépôt de résidus miniers</p>
<p>5 (1) Despite section 4, the owner or operator of a mine may deposit or permit the deposit of waste rock or an effluent that contains any concentration of a deleterious substance and that is of any pH into a tailings impoundment area that is either</p>	<p>5 (1) Malgré l'article 4, le propriétaire ou l'exploitant d'une mine peut rejeter — ou permettre que soient rejetés — des stériles ou un effluent, quel que soit le pH de l'effluent ou sa concentration en substances nocives, dans l'un ou l'autre des dépôts de résidus miniers suivants :</p>

Fig. 2. A fragment of the MMER

The MMER, which aims to protect aquatic life, is the Canadian regulation that directs metal mines to conduct *Environmental Effects Monitoring* (EEM) as a condition when depositing effluents resulting from mining activities [29]. The government supports this activity through an *EMM program* (EMMP). Consider a situation where *Environment and Climate Change Canada* (ECCC), the corresponding federal regulator, wants to review and report on the performance of MMER (a regulation) and of EEMP (a program), and interesting relationships between these two artefacts. With a scenario where four provincial metal mines (in Manitoba, Nova Scotia, Ontario, and Quebec) are reviewed between 2014 and 2016, we use GoRIM and Watson Analytics to describe how this can be achieved.

Building the Goal Models: To apply GoRIM, we first build goal models for MMER and EEMP using jUCMNav. An example for MMER, based on several fragments similar to the one shown in Fig. 2, is illustrated in Fig. 3. The structure of goal models enables us to capture the regulation/program structure (part/section, subpart/subsection, rule statements, etc.), and show different relationships (contributions and decompositions). A layer of indicators at the bottom enables the measurement of various aspects of rule statements, and some indicators can contribute to many rules. The model in Fig. 3 is a simplified version of MMER; the real GRL model for this regulation is much larger, but the selected subset is sufficient to illustrate GoRIM and investigate the functionalities of Watson Analytics. A similar model exists for the supporting program (EEMP, see Fig. 4).

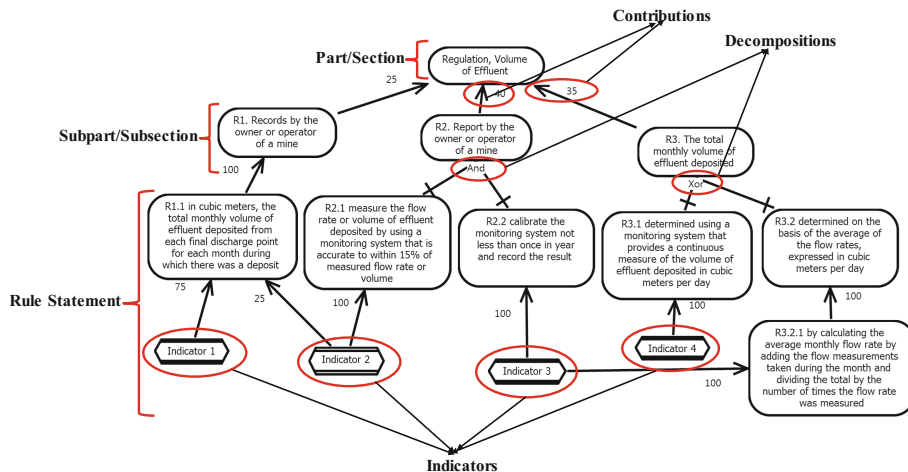


Fig. 3. Simplified MMER goal model

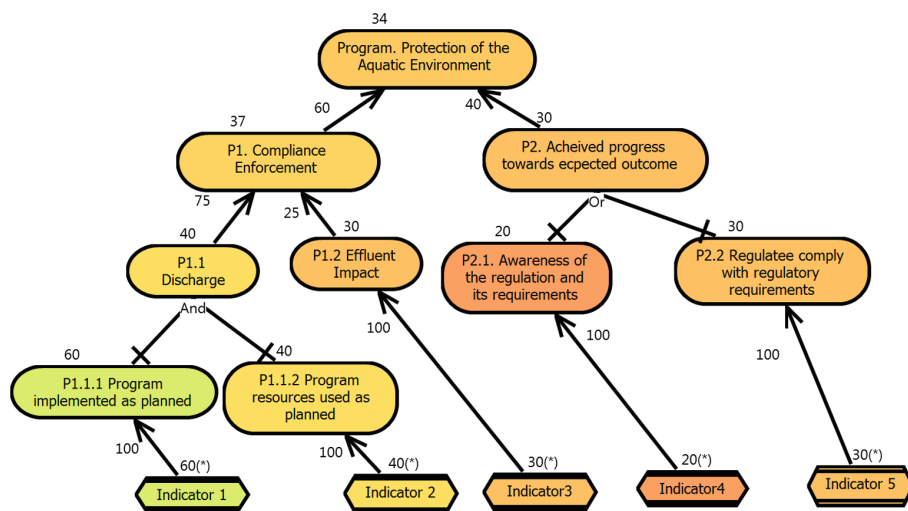


Fig. 4. Evaluated simplified EEMP goal model

Data Input, Evaluation, and Output: Next, we input data into the indicators of the regulation and program goal models via GRL strategies. We used sample compliance data for the four metal mines and evaluation data for EEMP activities from 2014 to 2016. The data we use here is synthetic as the real data and goal evaluations are confidential. GRL evaluation strategies representing each month of a year for each provincial mine were created. These GRL strategies define a set of initial values for the indicators of the MMER and EEMP goal models. Each indicator converts real

observation data (e.g., 35 parts per million) into a satisfaction level on a [0–100] scale through a comparison with target, threshold, and worst-case parameters.

We then used the GRL propagation algorithm to obtain compliance/performance satisfaction levels for the higher-level goals of the MMER and EEMP models. The sample EEMP goal model is shown with an evaluation for one of the mines at a specific time in Fig. 4. The computed satisfaction levels of the MMER and EEMP goal models, for different mines and months, are exported using jUCMNav’s export function, and can then be stored in a database for further analysis.

Data Visualization and Further Analysis: In our example, we exported the computed satisfaction levels as Comma Separated Value (CSV) files, imported as is into Watson Analytics. CSV files can be explored in Watson Analytics individually to provide more insight on the mines’ state of compliance with MMER and performance of EEMP between 2014 and 2016. However, much more value is offered when we combine both datasets. While this can be done using the query facilities in databases (if the data is stored in a database), in our example, we manually combined both CSV files, as illustrated in Fig. 5. The combination is based on shared dimensions of interest, for example, the months between 2014 and 2016 (*time*), the provinces where the mines exist (*location*), and the structure of the MMER and EEMP goal models (*regulation & regulatory program*). Each row in Fig. 5 stores the satisfaction value of each goal of the regulation and program models for a given provincial mine and a given month.

Location Dimension			Regulation & Regulatory Program Dimension																				
Month	Year	Province	MMER	R1_Owi	R2_Owi	R3_Toti	R1_1	MR2_1	VrR2_2	CrR3_1	DiR3_2	DiR3_2.1	EEM Program	P1_Conf	P2_Ach	P1.1	DI	P1.2	EF	P2.1_Rt	P2.2_Rt	P1.1.1	P1.1.2
January	2014	Manitoba	38	50	0	75	50	0	0	75	0	0	61	36	100	25	75	100	50	75	25	25	25
January	2015	Manitoba	80	100	50	100	100	75	50	100	50	50	49	49	50	50	50	50	50	50	50	50	50
January	2016	Manitoba	100	100	100	100	100	100	100	100	100	100	60	68	50	75	50	50	50	50	50	75	75
February	2014	Manitoba	33	25	25	50	25	50	25	50	25	25	40	0	100	0	0	100	50	25	0	25	0
February	2015	Manitoba	58	50	50	75	50	75	50	75	50	50	47	62	25	75	25	25	25	25	100	75	
February	2016	Manitoba	44	25	75	25	25	75	75	25	25	25	38	30	50	25	25	50	50	50	25	25	25

Fig. 5. Extract of the combined dataset for MMER and EEMP evaluations

Upon uploading the CSV file, Watson Analytics reviews the data and attempts to recognize automatically its nature. For example, it understands the meaning of the month, year, and province columns. It also uses multiple machine learning algorithms in parallel to analyze possible relationships in the data, leading to suggestions of questions and visualizations to be explored by the analyst. Figure 6 present six such *starting points*. Natural language questions in English can also be asked explicitly, as illustrated at the top of Fig. 6. When we asked the question “What is the relationship between MMER and EEM Program by year and province”, Watson Analytics analyzed our question (relationships in Watson are mainly quantitative correlations), automatically selected an appropriate visualization, and suggested further related questions based on our data. The suggested questions were sorted according to their computed relevance, as outlined in Table 1. These questions reflected correctly the time and location dimensions from the dataset. However, Watson Analytics does not understand the goal-oriented structure of regulations and programs, which is a

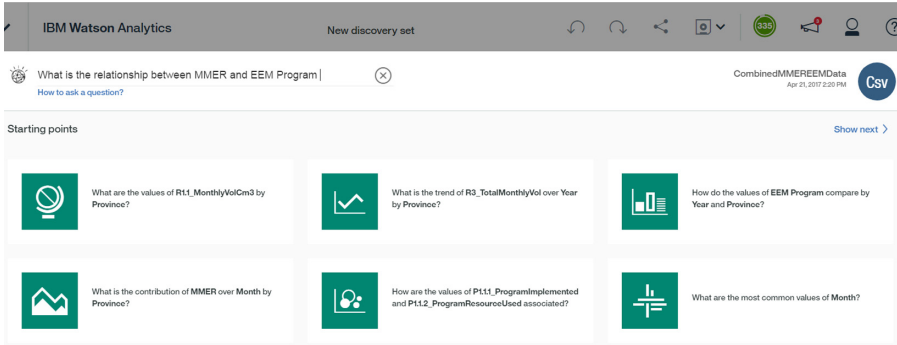


Fig. 6. Watson Analytics suggested questions, and interface to ask others based on the data

Table 1. Further questions suggested by Watson Analytics based on the initial question

Questions	Relevance Level
What is the relationship between the EEM Program and MMER by Province and Year?	Most Relevant
What are the values of EEM Program and MMER by Province across Year?	Most Relevant
What is the breakdown of EEM Program by Year and Province?	Most Relevant
What are the values of EEM Program and MMER for Year and Province?	Most Relevant
How do the values of EEM Program and MMER compare by Provinces across Year?	Most Relevant
How are the values of EEM Program and MMER associated across Year	Somewhat Relevant

domain-specific dimension. Although this limitation hurts the use of default navigation features at the user interface level (e.g., drilling up or down along this structure dimension), the aggregate satisfaction along such structure is still available as it was computed explicitly in the goal model by jUCMNav. For example, as shown in Fig. 4, the program satisfaction (34) combines the satisfactions of P1 (37) and P2 (30).

The Watson Analytics offering of related questions is necessary for exploring data as it provides opportunities to refine questions concerning insights desired from satisfaction data. Some proposed questions may actually have results supported by strong evidence, but still be irrelevant. Minimizing such noise still needs to be explored in future work.

Upon selecting the question “How do the values of EEM Program and MMER compare by Provinces across Year?”, Watson Analytics offered the visualization shown in Fig. 7. Such a visualization offers an opportunity to analyze MMER and the EEM program together along dimensions unavailable prior to this article. For example, the regulator (ECCC) can observe that while the performance of the EEM program (the bar charts) is fairly consistent across years on average, the yearly average compliance with MMER (the trend lines) has been inconsistent. From 2014 to 2016, the metal



Fig. 7. Visual comparison of the EEM program and MMR by province, across years

mines in the provinces of Manitoba and Quebec have had a growing increase in their compliance levels while there has been a decrease in Ontario for the same period. The metal mine in the province of Nova Scotia had an increase in its compliance level between 2014 and 2015 and a decrease in 2016 close to its level in 2014.

With this information, the ECCC can investigate possible explanations for these observations, including what drives the MMR (Fig. 8) and the EEM program (Fig. 9).

This analysis indicates that while “R3.2.1_AvgMonthlyFlowRate” (from the goal model in Fig. 3) and “Month” is the main driver of the MMR with 46%, while “P1_ComplianceEnforcement” (from the goal model in Fig. 4) and “Province” is the main driver for the EEM program with 57%. The ECCC could explore these candidate

What drives **MMER** ?

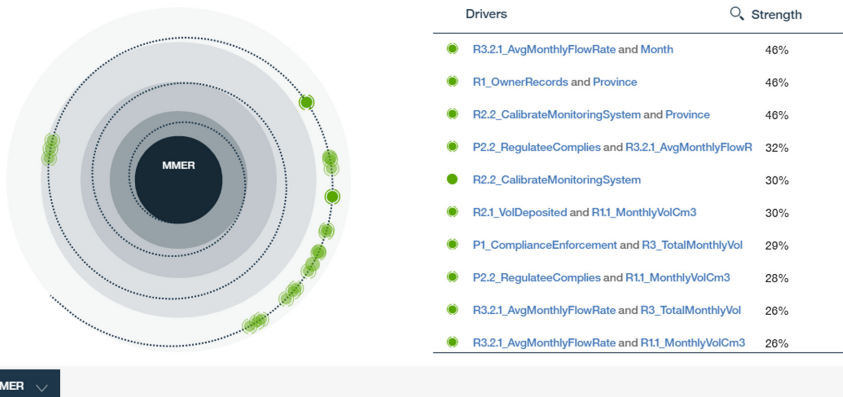


Fig. 8. Visualization showing what drives the MMR

What drives EEM Program ?

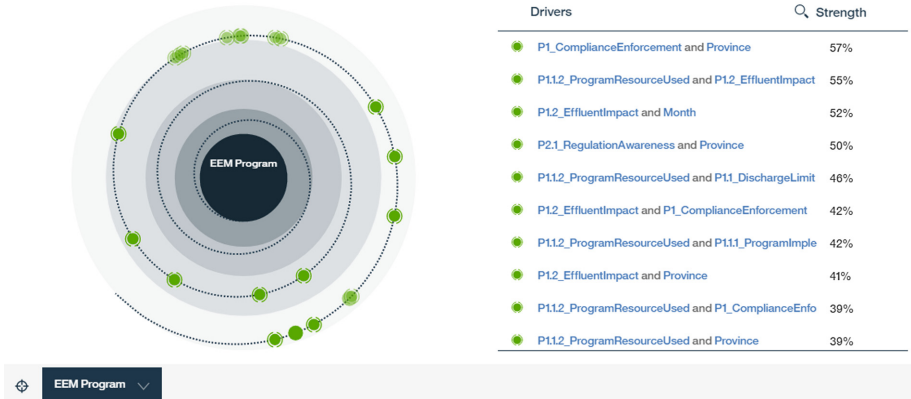


Fig. 9. Visualization showing what drives the EEM program

explanations even further or explore other visualizations or questions within Watson Analytics. Note that some proposed drivers have a low level of relevance in our context. For example, a contribution from R.3.2.1 to R3 is something known from the very structure of the goal model. How to prune out such known drivers in Watson is still under study.

6 Lessons Learned

Our exploration of GoRIM to analyze and report on compliance with a regulation (MMER) and the performance of a companion program (EEMP) were done by a Ph.D. student and two undergraduate students, under the supervision of two professors. All are co-authors of this article. The undergraduate students had no experience whatsoever with goal modelling or regulatory intelligence, whereas the PhD student and both professors had advanced knowledge. None of the participants had used Watson Analytics.

We have found Watson Analytics easy to learn and intuitive in visualizing and exploring compliance levels of regulations and performance levels of programs. While GRL offers appropriate tool-supported concepts for modeling regulations/programs and for analyzing them for one regulated party at one moment in time, Watson Analytics adds a form of analysis that was unavailable until now. Watson Analytics enables, out of the box, a combined analysis of data coming from multiple models (regulation and program), for many regulated parties and many moments in time, with opportunities for slicing, dicing, and drilling along several conventional dimensions (time and location). Suggestions for visualizations and related questions are also useful features.

During our exploration, we have learned the following lessons:

- (a) GoRIM: We can use a uniform goal-based modelling technique for regulations and regulatory programs. Opportunities exist to use this conceptualization to model more regulatory elements such as risks or intended societal outcomes.
- (b) GoRIM: Our method has the potential to accommodate the Big Data involved in the regulatory domain. Based on questions inquiring on how regulations and programs perform, indicators can be structured to collect data and feed goal models whose evaluations are further explored using data visualization tools.
- (c) Watson Analytics: We can ask many types of questions based on different dimensions in our data using natural language. This is in addition to the relevant “starting point” questions Watson Analytics provides upon loading data.
- (d) Watson Analytics: This tool is simple to learn and easy to use, requiring no data analytics, business intelligence, or specialized IT skills hence. This is a huge improvement upon the use of a conventional and heavyweight business intelligence tool (e.g., IBM Cognos) for regulatory intelligence, which was partially explored in [21]. This can lead to reductions in training costs and increased efficiency.
- (e) Watson Analytics: We also like the suggestions on how to visualize and analyze data through a variety of visualization alternatives that can be tailored to suit the types of questions being asked. This is in addition to user-defined functions, which can enable regulators to explore relationships not supported by default.

7 Limitations

Although we could create and analyze goal models of regulations (MMER) and programs (EEMP) using GoRIM, we are yet to evaluate GoRIM in a real-life context. We used fake data for our analysis and have not yet obtained feedback on the usefulness of GoRIM from the regulator who provided the example used here. There are also scalability and usability concerns to be addressed due to the size of real regulations. For example, the main section of the MMER goal model, excluding the eight different schedules that further explain the regulation, already contains 273 goals. We have not yet explored the usability of Watson Analytics in the presence of hundreds of goals, as well as the impact of changes to the goal models themselves (e.g., with the addition or deletion of goals, and hence of values in the database) on the analysis features of this tool. We are also yet to explore a more complete set of functionalities from Watson Analytics for exploring regulatory intelligence data. Finally, we have observed several limitations in that technology (e.g., Watson Analytics currently understands months but not quarters, and countries/provinces/states but not cities), and there might be other such limitations affecting regulations and programs in specific domains.

8 Conclusion and Future Work

To administer regulations effectively, a feedback loop involving data from and within the regulatory process is necessary. This information exhibits the properties of Big Data, creating the need for advanced tools and technologies to enable analysis and visualization providing the insight required to make informed decisions when administering regulations and their supporting programs. This paper proposed and illustrated the Goal-oriented Regulatory Intelligence Method (GoRIM), which uses the same conceptualization (goal models) to capture and analyze regulations and regulatory programs, and supports a robust analysis of compliance data. Moreover, the method exploits Watson Analytics to analyze and report on observed compliance levels, explore what they mean relative to the regulation, and determine how supporting regulatory programs contribute to observed compliance. As such, GoRIM can offer much value to regulators who want to assess the performance of their regulations and programs. The simplicity to learn and use Watson Analytics offers an attractive, pragmatic, out-of-the-box solution to support regulatory intelligence activities.

In the near future, we plan to use GoRIM to model real regulations and programs and use real data to explore concrete relationships in the regulatory process. In addition to addressing some of the limitations from the previous section, a usability study would also help determine the parts of GoRIM and of Watson Analytics that are really usable and of value to regulators.

Acknowledgements. This work was supported financially by the National Science and Engineering Research Council of Canada (NSERC) Discovery program. We are much thankful to Colette Lacroix and IBM Canada for access to Watson Analytics. We also thank Prof. Greg Richards, Dr. Randy Giffen, and Nick Cartwright for useful discussions, as well as the reviewers for their insightful suggestions.

References

1. OECD: Recommendation of the Council on Regulatory Policy and Governance. OECD Publishing, Paris (2012)
2. OECD: The Governance of Regulators, OECD Best Practice Principles for Regulatory Policy. OECD Publishing, Paris (2014)
3. Coglianese, C.: Measuring regulatory performance: evaluating the impact of regulation and regulatory policy. OECD Expert Paper No. 1, OECD Publishing, Paris (2012)
4. Nielsen, V.L., Parker, C.: Is it possible to measure compliance? Legal Studies Research Paper No. 192, Faculty of Law, The University of Melbourne (2006). SSRN <https://ssrn.com/abstract=935988>
5. Parker, D., Kirkpatrick, C.: Measuring regulatory performance. The economic impact of regulatory policy: a literature review of quantitative evidence. OECD Expert Paper No. 3, OECD Publishing, Paris (2012)
6. Horkoff, J., Aydemir, F.B., Cardoso, E., Li, T., Maté, A., Paja, E., Giorgini, P.: Goal-oriented requirements engineering: a systematic literature map. In: 24th International Requirements Engineering Conference (RE), pp. 106–115. IEEE CS (2016). doi:[10.1109/RE.2016.41](https://doi.org/10.1109/RE.2016.41)

7. Akhigbe, O., Amyot, D., Richards, G.: Information technology artifacts in the regulatory compliance of business processes: a meta-analysis. In: Benyoucef, M., Weiss, M., Mili, H. (eds.) *MCETECH 2015*. LNBIP, vol. 209, pp. 89–104. Springer, Cham (2015). doi:[10.1007/978-3-319-17957-5_6](https://doi.org/10.1007/978-3-319-17957-5_6)
8. Gandomi, A., Haider, M.: Beyond the hype: big data concepts, methods, and analytics. *Int. J. Inf. Manag.* **35**(2), 137–144 (2015). doi:[10.1016/j.ijinfomgt.2014.10.007](https://doi.org/10.1016/j.ijinfomgt.2014.10.007)
9. IBM: IBM Watson Analytics: Analytics Made Easy. <https://www.ibm.com/analytics/watson-analytics/us-en/index.html>. Accessed 23 Apr 2017
10. Akhigbe, O., Amyot, D., Mylopoulos, J., Richards, G.: What can information systems do for regulators? A review of the state-of-practice in Canada. In: *IEEE 11th International Conference on Research Challenges in Information Science (RCIS)*. IEEE CS (2017)
11. Siena, A., Perini, A., Susi, A., Mylopoulos, J.: A meta-model for modelling law-compliant requirements. In: *Requirements Engineering and Law (RELAW)*, pp. 45–51. IEEE CS (2009). doi:[10.1109/RELAW.2009.1](https://doi.org/10.1109/RELAW.2009.1)
12. Ingolfo, S., Siena, A., Perini A., Susi, A., Mylopoulos, J.: Modeling laws with Nòmos 2. In: *6th International Workshop on RE and LAW (RELAW)*, pp. 69–71. IEEE CS (2013). doi:[10.1109/RELAW.2013.6671350](https://doi.org/10.1109/RELAW.2013.6671350)
13. Ingolfo, S., Jureta, I., Siena, A., Perini, A., Susi, A.: Nòmos 3: legal compliance of roles and requirements. In: Yu, E., Dobbie, G., Jarke, M., Purao, S. (eds.) *ER 2014*. LNCS, vol. 8824, pp. 275–288. Springer, Cham (2014). doi:[10.1007/978-3-319-12206-9_22](https://doi.org/10.1007/978-3-319-12206-9_22)
14. Islam, S., Mouratidis, H., Jürjens, J.: A framework to support alignment of secure software engineering with legal regulations. *Softw. Syst. Model.* **10**(3), 369–394 (2011). doi:[10.1007/s10270-010-0154-z](https://doi.org/10.1007/s10270-010-0154-z)
15. Amyot, D., Mussbacher, G.: User requirements notation: the first ten years, the next ten years. *J. Softw. (JSW)* **6**(5), 47–768 (2011)
16. Ghanavati, S., Amyot, D., Peyton, L.: Compliance analysis based on a goal-oriented requirement language evaluation methodology. In: *17th IEEE International Requirements Engineering Conference (RE 2009)*, pp. 133–142. IEEE CS (2009). doi:[10.1109/RE.2009.42](https://doi.org/10.1109/RE.2009.42)
17. Ghanavati, S., Amyot, D., Rifaut, A.: Legal goal-oriented requirement language (Legal GRL) for modeling regulations. In: *6th International Workshop on Modeling in Software Engineering (MiSE)*, pp. 1–6. ACM (2014). doi:[10.1145/2593770.2593780](https://doi.org/10.1145/2593770.2593780)
18. Shamsaei, A., Pourshahid, A., Amyot, D.: Business process compliance tracking using key performance indicators. In: zur Muehlen, M., Su, J. (eds.) *BPM 2010*. LNBIP, vol. 66, pp. 73–84. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-20511-8_7](https://doi.org/10.1007/978-3-642-20511-8_7)
19. Amyot, D., et al.: Towards advanced goal model analysis with jUCMNav. In: Castano, S., Vassiliadis, P., Lakshmanan, L.V., Lee, M.L. (eds.) *ER 2012*. LNCS, vol. 7518, pp. 201–210. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-33999-8_25](https://doi.org/10.1007/978-3-642-33999-8_25)
20. Tawhid, R., Braun, E., et al.: Towards outcome-based regulatory compliance in aviation security. In: *20th IEEE International Requirements Engineering Conference (RE)*, pp. 267–272. IEEE CS (2012). doi:[10.1109/RE.2012.6345813](https://doi.org/10.1109/RE.2012.6345813)
21. Badreddin, O., Mussbacher, G., et al.: Regulation-based dimensional modeling for regulatory intelligence. In: *6th International Workshop on Requirements Engineering and Law (RELAW)*, pp. 1–10. IEEE CS (2013). doi:[10.1109/RELAW.2013.6671340](https://doi.org/10.1109/RELAW.2013.6671340)
22. Aggarwal, M., Madhukar, M.: IBM’s Watson Analytics for health care: a miracle made true. In: *Cloud Computing Systems and Applications in Healthcare*, pp. 117–134. IGI Global (2017). doi:[10.4018/978-1-5225-1002-4.ch007](https://doi.org/10.4018/978-1-5225-1002-4.ch007)
23. Anderson, F.: Watson Analytics sessions at World of Watson 2016. <https://www.ibm.com/communities/analytics/watson-analytics-blog/event-watson-analytics-sessions-at-world-of-watson-2016/>. Accessed 23 Apr 2017

24. Hynes, C.: Regulatory intelligence: implications for product development. In: 2014 TOPRA Module: Strategic Planning in Regulatory Affairs. <http://bit.ly/2pr5UiY>. Accessed 23 Apr 2017
25. Felgate, T.: What is regulatory intelligence? <http://www.regulatory-intelligence.eu/2013/02/what-is-regulatory-intelligence.html>. Accessed 23 Apr 2017
26. Maguire, P.: What is 'regulatory intelligence?'. Regulatory Affairs Professional Society. <http://bit.ly/2oWKRNe>. Accessed 23 Apr 2017
27. Rashidi-Tabrizi, R., Mussbacher, G., Amyot, D.: Transforming regulations into performance models in the context of reasoning for outcome-based compliance. In: 6th International Workshop on Requirements Engineering and Law (RELAW), pp. 34–43. IEEE CS (2013)
28. Pourshahid, A., Amyot, D., Peyton, L., Ghanavati, S., Chen, P., Weiss, M., Forster, A.J.: Business process management with the user requirements notation. *Electron. Commer. Res.* **9**(4), 269–316 (2009). doi:[10.1007/s10660-009-9039-z](https://doi.org/10.1007/s10660-009-9039-z)
29. Justice Laws Website: Consolidated federal laws of Canada, Metal Mining Effluent Regulations. <http://laws-lois.justice.gc.ca/eng/regulations/SOR-2002-222/>. Accessed 23 Apr 2017